

Distributed Approach for Deblurring Large Images with Shift-Variant Blur

Rahul Mourya*, André Ferrari†, Rémi Flamary†, Pascal Bianchi* and Cédric Richard†

*LTCI, Télécom ParisTech, Université Paris-Saclay, Paris 75013, France

†Laboratoire Lagrange, Observatoire de la Côte d'Azur, Université de Nice Sophia Antipolis, Nice 06108, France

Abstract—Image deblurring techniques are effective tools to obtain high quality image from acquired image degraded by blur and noise. In applications such as astronomy and satellite imaging, size of acquired images can be extremely large (up to gigapixels) covering a wide field-of-view suffering from shift-variant blur. Most of the existing deblurring techniques are designed to be cost effective on a centralized computing system having a shared memory and possibly multicore processor. The largest image they can handle is then conditioned by the memory capacity of the system. In this paper, we propose a distributed shift-variant image deblurring algorithm in which several connected processing units (each with reasonable computational resources) can deblur simultaneously different portions of a large image while maintaining a certain coherency among them to finally obtain a single crisp image. The proposed algorithm is based on a distributed Douglas-Rachford splitting algorithm with a specific structure of the penalty parameters used in the proximity operator. Numerical experiments show that the proposed algorithm produces images of similar quality as the existing centralized techniques while being distributed and being cost effective for extremely large images.

Index Terms—Distributed optimization, proximal projection, shift-variant blur, inverse problems, image deblurring.

I. INTRODUCTION

For many applications, it is essential to have high resolution images for precise analysis and inference. However, a certain amount of degradation (blur and noise) in acquired images is inherent to many imaging systems due to several physical factors. The image deblurring techniques are proven to be economical ways to enhance resolution, and signal-to-noise ratio. In general, image deblurring is an ill-posed inverse problem. In Bayesian setting, it is expressed as *maximum-a-posteriori* estimation problem, which can be casted to the following optimization problem

$$\arg \min_{\mathbf{x}} \{ \Psi_{\text{data}}(\mathbf{y}, \mathbf{x}) + \lambda \Psi_{\text{prior}}(\mathbf{x}) \} \quad (1)$$

where vectors $\mathbf{y} \in \mathbb{R}^m$ and $\mathbf{x} \in \mathbb{R}^n$ represent the two dimensional (2D) acquired (observed) and unknown (true) images, obtained by lexicographically ordering their pixels into vectors. Ψ_{data} is called data-fidelity term and depends upon the noise and the image formation model. Ψ_{prior} is called regularizer and promotes some prior knowledge on the unknown image \mathbf{x} . The scalar parameter λ keeps trade-off between these two terms.

Blur in an acquired image can be characterized by the point-spread-function (PSF). For a narrow field-of-view, the

blur is shift-invariant, and blurring is then a convolution between image and PSF. In the case of wide field-of-view, the blur is usually shift-variant, and then blurring is not a mere convolution. There does not exist any efficient and straight-forward way for shift-variant deblurring. However, there are some fast approximations proposed in [1,2]; see [3] for detailed comparison between the different approximations. If we approximate the noise in observed image by a non-stationary white Gaussian noise as considered in [4] (see [5] for more refined noise model), the discrete image formation model can be written as: $\mathbf{y} = \mathbf{H} \mathbf{x} + \varepsilon$. The matrix $\mathbf{H} \in \mathbb{R}^{m \times n}$ denotes the blurring operator. Usually $m < n$ since the observed image is restricted by the physical size of the sensor. ε denotes a zero-mean non-stationary white Gaussian noise, i.e. $\varepsilon(\ell) \sim \mathcal{N}(0, \sigma^2(\ell))$ with $\sigma^2(\ell)$ denoting the noise variance at the ℓ th entry of \mathbf{y} .

Let us introduce some more notations. Hereafter, upper-case and lower-case bold letters will denote matrices and vectors, respectively. If \mathcal{X} denotes Euclidean space, then $\langle \cdot, \cdot \rangle$ denotes the standard inner product on \mathcal{X} , and $\| \cdot \|$ represents Euclidean norm. If \mathbf{V} denotes a positive definite linear operator of \mathcal{X} onto itself, then $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbf{V}} = \langle \mathbf{x}, \mathbf{V} \mathbf{y} \rangle$, and $\| \cdot \|_{\mathbf{V}}$ represents the corresponding norm. When \mathbf{V} is a diagonal operator of the form $(\mathbf{V} \mathbf{x}) : \ell \mapsto \alpha(\ell) \mathbf{x}(\ell)$, we equivalently denote $\| \mathbf{x} \|_{\mathbf{V}}^2$ by $\| \mathbf{x} \|_{\alpha}^2 = \sum_{\ell} \alpha(\ell) \mathbf{x}(\ell)^2$. If \mathbf{L} is a linear operator, then \mathbf{L}^{T} denotes the adjoint operator.

Considering the above image formation model, the image deblurring problem (1) can be explicitly expressed as

$$\arg \min_{\mathbf{x} \geq 0} \left\{ \frac{1}{2} \| \mathbf{y} - \mathbf{H} \mathbf{x} \|_{\mathbf{W}}^2 + \lambda \phi(\mathbf{D} \mathbf{x}) \right\} \quad (2)$$

where the matrix \mathbf{W} is diagonal with its elements given by $\mathbf{W}(\ell, \ell) = 1/\sigma^2(\ell)$ for the observed pixels, and $\mathbf{W}(\ell, \ell) = 0$ for unmeasured pixels. The function ϕ is some regularizer and \mathbf{D} some linear operator, e.g., finite forward difference or some discrete wavelet transform. Depending upon the structure of the two terms in (2), the solution can be found by any fast algorithm from the existing vast literature on numerical optimization methods. We refer the readers to [6]–[9] and the references therein for the recent and fast optimization algorithms, especially for inverse problems in imaging.

Motivation: Image deblurring is a well studied topic, and there is a vast literature [4,5,7,8,10] for moderate size (few megapixels) images with shift-invariant blur. The applications in astronomy, satellite imagery and others are able to capture

The work is supported by the Agence Nationale pour la Recherche (ANR), France, (ODISSEE project, ANR-13-ASTR-0030 and MAGELLAN project, ANR-14-CE23-0004-01).

extremely large images (up to gigapixels) of wide field-of-views suffering from shift-variant blur. Image deblurring becomes further complicated when shift-variant blur occurs. Some works [1]–[3] proposed efficient shift-variant image deblurring methods. However, all the methods listed above are designed and implemented to be efficient on a centralized computing system having a shared memory space, and possibly multicore processor. Hereafter, we will refer to them as *centralized* deblurring methods. The largest image they can deblur is limited by the capacity of physical memory available on the system. Practically, it is not possible to build a single centralized system having several processor cores and a huge shared memory. Modern distributed computing infrastructures are proving to be far more cost effective for solving huge-scale problems than centralized systems. A distributed computing system consists of several connected processing nodes each having a reasonable number of processors and sufficient memory capacity. The nodes communicate together to achieve a common goal. Many applications in machine learning, data mining, and others are already benefiting from distributed computing approaches for learning patterns from huge datasets possibly located at different locations. As per our knowledge, there are very few works [11,12] considering distributed computing approaches for image restoration/reconstruction.

Contribution: In this paper, we propose a distributed approach for deblurring large images suffering from shift-variant blur. Our approach is based on splitting the large acquired image into sufficiently small overlapping blocks provided that one is able to sample local PSFs within each block. We reformulate the original problem (2) into a distributed optimization problem by introducing certain approximations. Then, we solve it with a distributed version of Douglas-Rachford (D-R) splitting algorithm by imposing certain consensus among the blocks such that we finally obtain a deblurred image without discontinuity among the blocks. The experimental results show that our distributed image deblurring method is able to achieve similar quality of deblurred image as *centralized* deblurring methods while allowing distribution, and thus being cost effective for extremely large image.

II. DISTRIBUTED APPROACH FOR IMAGE DEBLURRING

High-speed links among the nodes of distributed system may not be always possible. Thus, an efficient distributed algorithm is the one which is computation intensive rather than communication intensive. For shift-variant image deblurring, a possible approach is to use the distributed array abstraction available on modern distributed system, and reimplement the existing *centralized* methods [2,3] using distributed arrays. However, the bottleneck of such approaches is extensive data communication among the nodes at each iteration of the optimization algorithm, e.g., each time \mathbf{H} or \mathbf{H}^T is applied, the nodes need to exchange a large amount of data. To avoid such a bottleneck, a possible approach is to split the large image into sufficiently small contiguous blocks, and deblur them locally at the different nodes while exchanging information to maintain coherencies among the blocks, but only after certain

intervals or certain suboptimality is met. The fast approximation of shift-variant blur operator based on PSF interpolation described in [3] provides an interesting idea in this direction. If blur varies smoothly in field-of-view, then PSF at any location in the field can be well approximated by linear combination (interpolation) of only few PSFs sampled on a regular grid points within the field. First-order 2D interpolation weights are sufficient to render smooth variation of the blur within the field while keeping the shift-variant blur operator computationally efficient [2,3]. This suggests us to split the acquired image into overlapping blocks, approximate the regularizer on the whole image by sum of the regularizers on the blocks, and then deblur the blocks locally while imposing consensus on overlapping pixels by weighted average using first-order 2D interpolation weights. This will result into a deblurred image without any discontinuities among the blocks. Moreover, any incoherencies arising due to the approximation made on the regularizer will be compensated by the overlaps and the consensus among them. Below, we present formally the proposed distributed approach for shift-variant deblurring.

A. General Setting

Consider a distributed computing system with a set of N connected processing nodes. Given an observed image $\mathbf{y} \in \mathbb{R}^m$, and the PSFs $\mathbf{h}_i, i = 1, \dots, N$ sampled at regular grid points within the field-of-view, we split \mathbf{y} into overlapping blocks $\mathbf{y}_i \in \mathbb{R}^{m_i}, i = 1, \dots, N$. As in [3], we consider first-order 2D interpolation weights $\omega_i, i = 1, \dots, N$ corresponding to the overlapping blocks \mathbf{y}_i . Fig. 1 illustrates the splitting, and the shape of the interpolation weights for 1D signal. We distribute the blocks \mathbf{y}_i , the corresponding local blur operators $\mathbf{H}_i \in \mathbb{R}^{m_i \times n_i}$, and the interpolation weights ω_i among the nodes. We then, seek to distributively estimate the whole unknown crisp image $\mathbf{x} \in \mathbb{R}^n$. Let us denote by $\mathcal{P}_1, \dots, \mathcal{P}_N$ a collection of N subsets of $\{1, \dots, n\}$. For every $i = 1, \dots, N$, we assume that the i th node is in charge of estimating the components of \mathbf{x} corresponding to the indices \mathcal{P}_i . The subsets $\mathcal{P}_1, \dots, \mathcal{P}_N$ are overlapping. Hence, different nodes handling a common component of the parameter vector \mathbf{x} must eventually agree on the value of the latter. We introduce formally the product space $\mathcal{X} := \mathbb{R}^{\mathcal{P}_1} \times \dots \times \mathbb{R}^{\mathcal{P}_N}$, and denote by \mathcal{C} the set of vectors $(\mathbf{x}_1, \dots, \mathbf{x}_N) \in \mathcal{X}$ satisfying the restricted consensus condition $\forall (i, j) \in \{1, \dots, N\}^2, \forall \ell \in \mathcal{P}_i \cap \mathcal{P}_j, \mathbf{x}_i(\ell) = \mathbf{x}_j(\ell)$. Moreover, we assume that every i th node is provided with a local convex, proper and lower semicontinuous function $f_i : \mathbb{R}^{\mathcal{P}_i} \rightarrow (-\infty, +\infty]$. We consider the following constrained minimization problem on $\mathbb{R}^{\mathcal{P}_1} \times \dots \times \mathbb{R}^{\mathcal{P}_N}$:

$$\arg \min_{\mathbf{x}_1, \dots, \mathbf{x}_N} \sum_{i=1}^N f_i(\mathbf{x}_i) \text{ s.t. } (\mathbf{x}_1, \dots, \mathbf{x}_N) \in \mathcal{C}. \quad (3)$$

The local cost function f_i is composed of the local data-fidelity term $\frac{1}{2} \|\mathbf{y}_i - \mathbf{H}_i \mathbf{x}_i\|_{\mathbf{W}_i}^2$ for some positive semidefinite \mathbf{W}_i as in (2), a regularizer ϕ_i and positivity constraint on \mathbf{x}_i , i.e., $\mathbf{x}_i \in \mathbb{R}_+^{\mathcal{P}_i}$. If \mathcal{A} is a set, the notation $\iota_{\mathcal{A}}$ stands for the indicator function of the set \mathcal{A} , equal to zero on that set and




Fig. 1. An 1D illustration showing the splitting of the observed image, the extent of the overlap, and shape of the interpolation weights. The crossbar on \mathbf{y} are the regular grid points where \mathbf{h}_i are sampled. \mathbf{x}_i are the locally deblurred blocks at i th nodes obtained from the corresponding \mathbf{y}_i . The dashed parts at the ends of each \mathbf{x}_i are the extra pixels estimated at boundaries assuming no measurements were available for them. ω_i are the interpolation weights corresponding to the support of \mathbf{x}_i , and its values are within range $[0, 1]$ such that $\sum_{i=1}^3 \omega_i(\ell) = 1, \forall \ell = 1, \dots, n$.

to $+\infty$ elsewhere. Thus, the local cost functions needed to be minimized at the nodes have the form:

$$f_i(\mathbf{x}_i) = \frac{1}{2} \|\mathbf{y}_i - \mathbf{H}_i \mathbf{x}_i\|_{\mathbf{W}_i}^2 + \lambda \phi_i(\mathbf{D}_i \mathbf{x}_i) + \iota_{\mathbb{R}_+^{\mathcal{P}_i}}(\mathbf{x}_i)$$

where ϕ_i are convex, proper and lower semicontinuous functions and \mathbf{D}_i are linear operators on $\mathbb{R}^{\mathcal{P}_i}$. Thus, the overall cost functions $f_i, i = 1, \dots, N$ are convex, proper and lower semicontinuous functions, and so is the problem (3). Note that \mathbf{x}_i are estimated little larger than \mathbf{y}_i to avoid boundary artifact arising due to circular blur assumption; see [8] for details.

B. Optimization Algorithm

For any convex, proper and lower semicontinuous function $h : \mathcal{X} \rightarrow (-\infty, +\infty]$, let us introduce the *proximity operator* $\text{prox}_{\mathbf{V}^{-1}, h}(\mathbf{v}) = \arg \min_{\mathbf{w} \in \mathcal{X}} h(\mathbf{w}) + \frac{\|\mathbf{w} - \mathbf{v}\|_{\mathbf{V}}^2}{2}$, for all $\mathbf{v} \in \mathcal{X}$. To solve (3), we consider D-R Splitting algorithm [13], and reformulate (3) as

$$\arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) + g(\mathbf{x}) \quad (4)$$

where $g = \iota_{\mathcal{C}}$ is the indicator function of \mathcal{C} and $f(\mathbf{x}) = \sum_i f_i(\mathbf{x}_i)$ for every $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ in \mathcal{X} . Let $\rho^{(k)}$ be a sequence in $]0, 2[$, and $\epsilon_j^{(k)}, j = 1, 2$ be sequences in \mathcal{X} , then the iterations of the D-R splitting algorithm writes as:

$$\begin{aligned} \mathbf{x}^{(k+1)} &= \text{prox}_{\mathbf{V}^{-1}, f}(\mathbf{u}^{(k)}) + \epsilon_1^{(k)} \\ \mathbf{z}^{(k+1)} &= \text{prox}_{\mathbf{V}^{-1}, g}(2\mathbf{x}^{(k+1)} - \mathbf{u}^{(k)}) + \epsilon_2^{(k)} \\ \mathbf{u}^{(k+1)} &= \mathbf{u}^{(k)} + \rho^{(k)} (\mathbf{z}^{(k+1)} - \mathbf{x}^{(k+1)}) \end{aligned}$$

Given that the set of minimizers of (3) is non-empty, and certain conditions on the sequences $\rho^{(k)}$ and $\epsilon_j^{(k)}, j = 1, 2$ are met, then the sequence $\mathbf{x}^{(k)}$ converges to a minimizer of (4) as $k \rightarrow \infty$; see [14, Corollary 5.2] for the proof. The parameter $\rho^{(k)}$ is referred to as relaxation factor that can be tuned to improve the convergence speed. The sequences $\epsilon_j^{(k)}, j = 1, 2$ allow some perturbations in the two prox operations, which is

very useful in the cases when the prox operations do not have closed-form solutions, and rely on some iterative solvers.

From now onward, we assume that \mathbf{V} is a diagonal operator of the form $\mathbf{V}\mathbf{x} = (\mathbf{V}_1\mathbf{x}_1, \dots, \mathbf{V}_N\mathbf{x}_N)$, where for every i ,

$$\begin{aligned} \mathbf{V}_i \mathbf{x}_i : \mathcal{P}_i &\rightarrow \mathbb{R} \\ \ell &\mapsto \alpha_i(\ell) \mathbf{x}_i(\ell), \end{aligned}$$

where $\alpha_i(\ell)$ is a positive coefficient specified later. For every $\ell \in \{1, \dots, n\}$, we introduce the set $\mathcal{P}_\ell^- = \{i : \ell \in \mathcal{P}_i\}$.

Lemma II.1. For every $\mathbf{x} \in \mathcal{X}$, the quantity $\mathbf{z} = \text{prox}_{\mathbf{V}^{-1}, g}(\mathbf{x})$ is such that for every $i \in \{1, \dots, N\}$ and every $\ell \in \mathcal{P}_i$,

$$\mathbf{z}_i(\ell) = \frac{\sum_{i \in \mathcal{P}_\ell^-} \alpha_i(\ell) \mathbf{x}_i(\ell)}{\sum_{i \in \mathcal{P}_\ell^-} \alpha_i(\ell)}. \quad (5)$$

Proof. The proof is straight-forward. The values of $\mathbf{z}_i(\ell)$ corresponding to the common components (overlapping pixels) among $\mathbf{x}_i, i = 1, \dots, N$ is given by weighted average of those components at the different nodes. \square

Using the above lemma, we present the proposed distributed image deblurring approach in Algorithm 1, which will be referred to as *proposed* deblurring method. The *proposed* method is a synchronous distributed optimization algorithm without any explicit master node. Given initial guesses of the local solutions at each node, the first step of Algorithm 1 executes the LOCAL-SOLVER simultaneously at all the nodes to obtain locally deblurred blocks, and then synchronizes all the nodes. In the second step, it exchanges the common components among the nodes, and performs the consensus operation distributively on all the nodes. Finally, it updates distributively the local copies of the variable $\mathbf{u}^{(k)}$.

Similar to the variants of ADMM [7]–[9], the convergence speed of D-R splitting algorithm is dependent upon the penalty parameters α_i . Selecting optimal values of penalty parameters for fast convergence is still an open challenge. We select $\alpha_i = \gamma \omega_i$ for $\gamma > 0$, so that we can tune γ to achieve possibly fast convergence, and as well impose smooth variation of the blur among adjacent blocks.

III. NUMERICAL EXPERIMENTS AND RESULTS

A. Experimental Setup

To evaluate the *proposed* deblurring method, we performed numerical experiment on shift-variant image deblurring under different parameter settings. We considered grayscale “Barbara” image, resized it to 1151×1407 pixels, and extended its dynamic range linearly to have maximum intensity up to 6000 photons/pixels. We will refer to this image as reference image shown in Fig. 3(b). We generated 9×9 grid of normalized Gaussian PSFs each of 201×201 pixels with the central PSF having full-width-half-maximum (FWHM) equal to 3.5×3.5 pixels, and increased the FWHM linearly along the radial direction up to 16.5×10.5 pixels for the PSF at extreme corner of the reference image as shown in Fig. 3(a). We blurred the reference image with these PSFs using shift-variant blur operator based on PSF interpolation [3]. We obtained the observed image, shown in Fig. 3(c), by adding white

Algorithm 1 PROPOSED DISTRIBUTED IMAGE DEBLURRING

```

procedure DISTRIBUTED-SOLVER
  Initialize:  $\mathbf{u}_i \leftarrow \mathbf{u}_i^{(0)}, \forall i = 1, 2, \dots, N$ 
  while not converged do
    for  $i = 1 \dots N$  do
       $\mathbf{x}_i \leftarrow \text{LOCAL-SOLVER}(\mathbf{u}_i; \alpha_i, f_i)$ 
    end for
    for  $\ell = 1 \dots n$  do
      Compute distributively at nodes  $i \in \mathcal{P}_\ell^-$ :
       $\bar{\mathbf{z}}_i(\ell) \leftarrow \sum_{i \in \mathcal{P}_\ell^-} \omega_i(\ell)(2\mathbf{x}_i(\ell) - \mathbf{u}_i(\ell))$ 
    end for
    for  $i = 1 \dots N$  do
       $\mathbf{u}_i(\ell) \leftarrow \mathbf{u}_i(\ell) + \rho(\bar{\mathbf{z}}_i(\ell) - \mathbf{x}_i(\ell))$  for all  $\ell \in \mathcal{P}_i$ 
    end for
  end while
  return  $\mathbf{x}_1, \dots, \mathbf{x}_N$ 
end procedure
procedure LOCAL-SOLVER( $\mathbf{u}; \alpha, f$ )
   $\mathbf{w} \leftarrow \text{prox}_{\alpha^{-1}, f}(\mathbf{u}) = \arg \min_{\mathbf{w}} \{f(\mathbf{w}) + \frac{1}{2}\|\mathbf{w} - \mathbf{u}\|_\alpha^2\}$ 
  return  $\mathbf{w}$ 
end procedure

```


(a) SNR vs λ

Fig. 2. Comparison of the three approaches in term of image quality (SNR) at different levels of regularization. The legends “Central-3x3”, “Proposed-3x3”, “Indpndt-3x3” denote the *centralized*, *proposed*, and *independent* methods, respectively, when using 3×3 grid of PSFs. Other legends have similar meanings. The dark line denotes the SNR of the observed image.

Gaussian noise of variance $\sigma^2 = 400$ to the blurred image. We deliberately chose low level noise in the observed image so that any incoherency arising in the deblurred image due to the approximations we made was not superseded by the strong regularization level required for low signal-to-noise ratio.

In order to compare the performance of the *proposed* method, we considered two other possible approaches. Since, the image we considered was not very large, thus the first method we selected was the *centralized* method, which solves the original problem (2) using the shift-variant blur operator based on PSF interpolation. The second method we selected was a naïve “split, deblur independently and blend” approach in which the final deblurred image was obtained by blending together the independently deblurred blocks. Blending is done using weighted averaging (5) with the same first-order 2D interpolation weights ω_i as in the *proposed* method. Though the second method is not intended to solve either the original problem (2) or the distributed problem (3), it is a straight-

forward way to deblur a large image with minimal computational resources. Hereafter, we will refer to the second method as *independent* deblurring method.

We chose ϕ_i be Huber loss function and \mathbf{D}_i be circular forward finite difference operator, so the regularizer writes as

$$\phi_i(\mathbf{D}_i \mathbf{x}_i) = \begin{cases} \frac{1}{2} \|\mathbf{D}_i \mathbf{x}_i\|_2^2, & \text{if } \|\mathbf{D}_i \mathbf{x}_i\|_2 \leq \delta \\ \delta(\|\mathbf{D}_i \mathbf{x}_i\|_2 - \frac{\delta}{2}), & \text{otherwise} \end{cases}$$

We chose this regularizer so that the local cost functions f_i is differentiable, and chose a quasi-Newton method: limited memory variable metric with bound-constraint (VMLM-B) [15] as the LOCAL-SOLVER. VMLM-B does not require any manual parameter tuning for fast convergence. Thus, this left us with only a single parameter γ to be tuned for fast convergence of the *proposed* method. Moreover, we used VMLM-B for solving the optimization in both the *centralized* and *independent* deblurring methods.

For all experiments, we fixed $\rho^{(k)} = 1$, and after few trials, we found that $\gamma = 0.001$ resulted into a fast convergence of Algorithm 1. All the results presented below were obtained after 25 iterations of Algorithm 1; it was observed that 25-30 iterations were generally sufficient. For the *centralized* and *independent* deblurring methods, we allowed VMLM-B to perform 1000 iterations or until it met its own stopping criterion.

B. Results

We considered two image quality metrics: signal-to-noise ratio (SNR) and Structural Similarity Index (SSIM) for comparing the quality of the deblurred images. We fixed heuristically the parameter $\delta = 100$, and performed the image deblurring for different values of λ in a wide range to see if one of the methods performs better than the others for certain values of λ . Moreover, we considered 3×3 , 5×5 and 6×6 grid of PSFs to see the effect of better approximation of shift-variant blur. The second row of Fig. 3 shows the deblurred images obtained by the three approaches when using 6×6 grid of PSFs. The plots in Fig. 2 compare the performance of the three deblurring approaches in term of image quality obtained for different values of λ , and different density of PSF grid. We observe that the naïve *independent* deblurring method, as expected, performed worst among the three. As pointed out above, this is due to the fact that it is the crudest and computationally cheapest way to perform shift-variant image deblurring by splitting image into pieces. The next important aspect we noticed is that the quality of deblurred image improves drastically when using denser grid of PSFs. This can be accounted by the fact that finer grid of PSFs are able to approximate more accurately the true variation of blur than a coarser grid of PSFs. A similar observation has been also noted in [3]. We also noticed that the quality of the deblurred image obtained by *proposed* method is slightly lower than those by *centralized* method. This could be due to two reasons: (i) some information at the boundaries of each deblurred blocks are lost, and (ii) the approximation made in the regularization can impact negatively for some image.

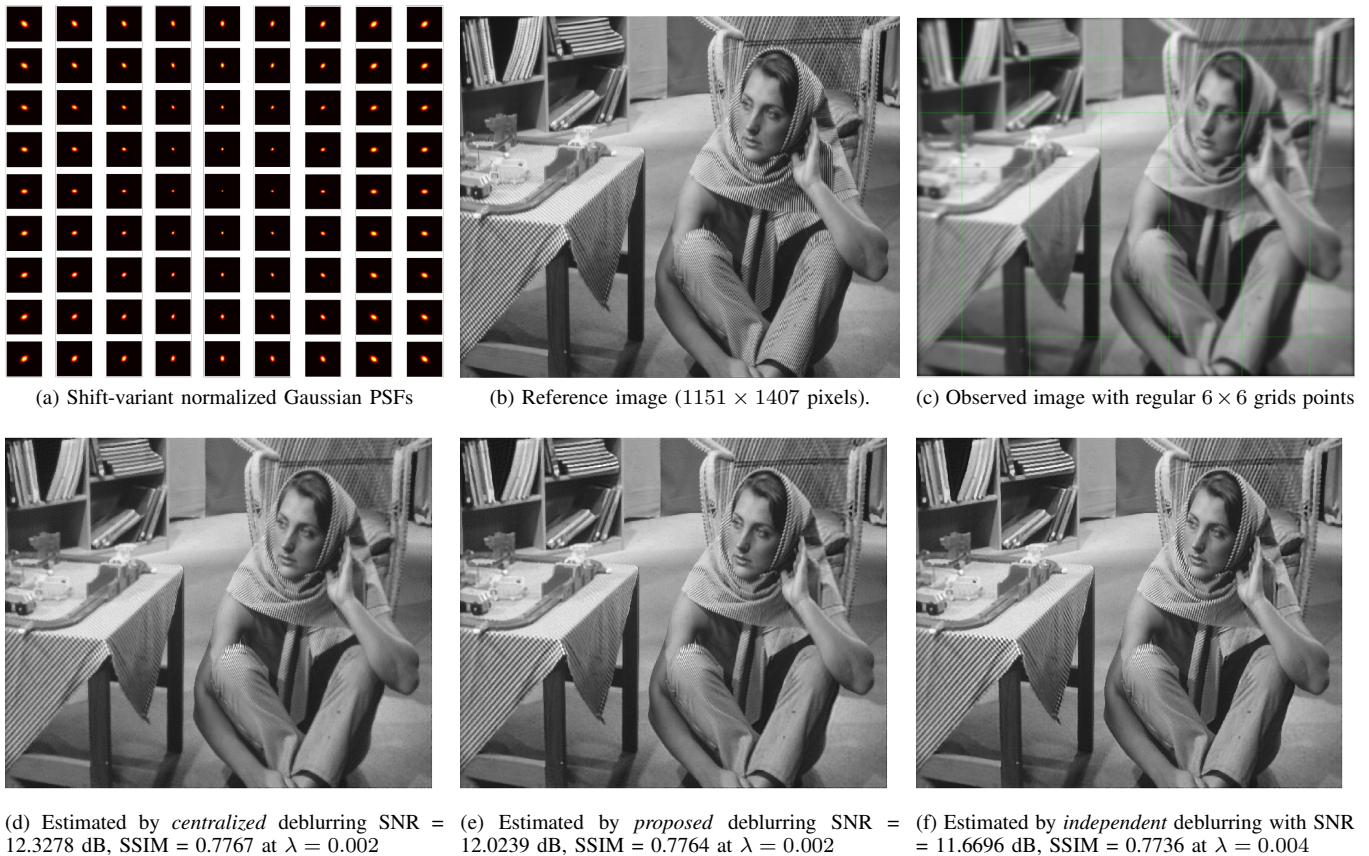


Fig. 3. Experimental setup for shift-variant deblurring and the deblurred images obtained by three methods using 6 × 6 grid of PSFs. The crossing points of the green lines in (c) are the grid points where PSFs are sampled.

IV. CONCLUSION

The paper proposed a distributed shift-variant deblurring algorithm scalable up to extremely large images that cannot be handled by the existing centralized deblurring methods. The proposed algorithm is rather generic in the sense that it can be easily extended to different applications, e.g., shift-invariant deblurring, or for different combination of data-fidelity and regularizer. The proposed algorithm is based upon distributed formulation of classical image deblurring problem by introducing certain approximations, which may slightly compromise the quality of the deblurred image. But, the compromise in the quality is highly compensated by the computational advantages of distributed computing.

REFERENCES

- [1] J. G. Nagy and D. P. O’Leary, “Restoring images degraded by spatially variant blur,” *SIAM Journal on Scientific Computing*, vol. 19, no. 4, pp. 1063–1082, 1998.
- [2] M. Hirsch, S. Sra, B. Sch, S. Harmeling, and B. Cybernetics, “Efficient Filter Flow for Space-Variant Multiframe Blind Deconvolution,” in *CVPR*, San Francisco, CA, 2010, pp. 607–614.
- [3] L. Denis, E. Thiébaud, F. Soulez, J.-M. Becker, and R. Mourya, “Fast Approximations of Shift-Variant Blur,” *International Journal of Computer Vision*, vol. 115, no. 3, pp. 253–278, 2015.
- [4] L. M. Mugnier, T. Fusco, and J.-M. Conan, “Mistral: a myopic edge-preserving image restoration method, with application to astronomical adaptive-optics-corrected long-exposure images,” *JOSA A*, vol. 21, no. 10, pp. 1841–1854, 2004.
- [5] F. Benvenuto, A. L. Camera, C. Theys, a. Ferrari, H. Lantéri, and M. Bertero, “The study of an iterative method for the reconstruction of images corrupted by Poisson and Gaussian noise,” *Inverse Problems*, vol. 24, no. 3, p. 035016, jun 2008.
- [6] A. Beck and M. Teboulle, “Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems,” *IEEE transactions on image processing*, vol. 18, no. 11, pp. 2419–34, 2009.
- [7] M. V. Afonso, J. M. Bioucas-dias, and M. A. T. Figueiredo, “An Augmented Lagrangian Approach to the Constrained Optimization Formulation of Imaging Inverse Problems,” *IEEE Transactions on Image Processing*, vol. 20, no. 3, pp. 681–695, 2011.
- [8] A. Matakos, S. Ramani, and J. A. Fessler, “Accelerated edge-preserving image restoration without boundary artifacts,” *IEEE Transactions on Image Processing*, vol. 22, no. 5, pp. 2019–2029, 2013.
- [9] R. Mourya, L. Denis, J.-M. Becker, and E. Thiébaud, “Augmented lagrangian without alternating directions: Practical algorithms for inverse problems in imaging,” in *IEEE ICIP*, Quebec, 2015, pp. 1205–1209.
- [10] W. H. Richardson, “Bayesian-based iterative method of image restoration,” *JOSA*, vol. 62, no. 1, pp. 55–59, 1972.
- [11] C. Meillier, P. Bianchi, and W. Hachem, “Two distributed algorithms for the deconvolution of large radio-interferometric multispectral images,” in *EUSIPCO*, Budapest, Hungary, 2016, pp. 728–732.
- [12] A. Ferrari, D. Mary, R. Flamary, and C. Richard, “Distributed image reconstruction for very large arrays in radio astronomy,” in *2014 IEEE 8th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, 2014, pp. 389–392.
- [13] P. L. Lions and B. Mercier, “Splitting Algorithms for the Sum of Two Nonlinear Operators,” *SIAM Journal of Numerical Analysis*, vol. 16, no. 6, pp. 964–979, 1979.
- [14] P. L. Combettes, “Solving monotone inclusions via compositions of nonexpansive averaged operators,” *Optimization*, vol. 53, no. 5-6, pp. 475–504, 2004.
- [15] E. Thiébaud, “Optimization issues in blind deconvolution algorithms,” in *Astronomical Telescopes and Instrumentation*. International Society for Optics and Photonics, 2002, pp. 174–183.